Speaker companion guide

Managing an R&D; portfolio through the critical value creation period (CVCP)

Companion guide to the ASHG 2026 talk: R&D; principles, the CVCP, human genetics, and AI

Speaker	Robert Plenge, MD, PhD, EVP and Chief Research Officer, Bristol Myers Squibb
Session	Translating Human Genetics into Better Drug Development Decisions
Date	Tuesday 20 October 2026, 13:00-15:00, Montréal
Follows	Jim, Sally and Mark; precedes the panel with Beryl Cummings and Pradeep Natarajan

Executive summary

This guide accompanies Robert Plenge's talk, *Managing an R&D portfolio through the critical value creation period (CVCP)*, at the 5 Prime Sciences Symposium at ASHG 2026 (Montréal, 20 October 2026). It follows the slides in order, adds external context, and highlights where human genetics changes the decision.

The central argument is that the most value in drug R&D is created on the path to clinical proof of concept (PoC), the critical value creation period, and that a portfolio should be managed by how efficiently each program buys a change in probability of success (PoS) during that period.

- **The problem and the solutions.** Biopharma R&D is failure-prone, expensive and inefficient. The solutions are rigorous R&D principles combined with managing through clinical PoC, and integrating AI and automation.
- **The CVCP.** The most value is created on the path to clinical PoC, the first point at which a stable product can be told apart from a mega-blockbuster.
- **Four actions.** Large companies need some mega-blockbusters (over USD 5B in annual revenue) to grow. Four actions follow: minimize time and capital to PoC; increase PoS at PoC and approval; increase the ratio of mega-blockbusters to stable products; and increase speed to market after the CVCP.
- **Managing inflection points.** R&D principles set the starting PoS; the key measure is dollars spent per change in PoS; and revenue potential, set by the aspirational target product profile (aTPP), is refined through PoC.
- **Three programs.** An HbF-inducing degrader for sickle cell disease, a BCL6 degrader for lymphoma and CD19 CAR T cells for autoimmune disease draw their causal evidence from germline genetics, somatic genetics and human pharmacology respectively. An illustrative portfolio map plots conviction against CVCP value.
- **AI.** BMS is investing toward an AI-enabled research organization (the dream) and building the capabilities to realize it (the build). External benchmarks such as DrugTargetWorld show that AI agents can run causal-genetics analyses but do not yet reliably make the integrative judgments decisions require.

Human genetics runs through every step: it raises the starting PoS, supplies the biomarkers and patient selection that make each inflection point informative, and informs the breadth of the aTPP. A linked reference list supports further reading.

1. Problem statements and solutions

Biopharma R&D is **failure-prone, expensive and inefficient** (Figure 1). These three problems frame the session. Earlier speakers described how genetic evidence is interpreted, translated and turned into conviction. This talk addresses how that evidence should be used to manage investment.

Two solutions address all three problems:

1. **Adhering to rigorous R&D principles and managing through clinical PoC.** This is the subject of Sections 2 to 7.
2. **Integrating AI and automation** to improve efficiency, shorten decision cycles and enhance the quality of IND output. This is the subject of Section 8.

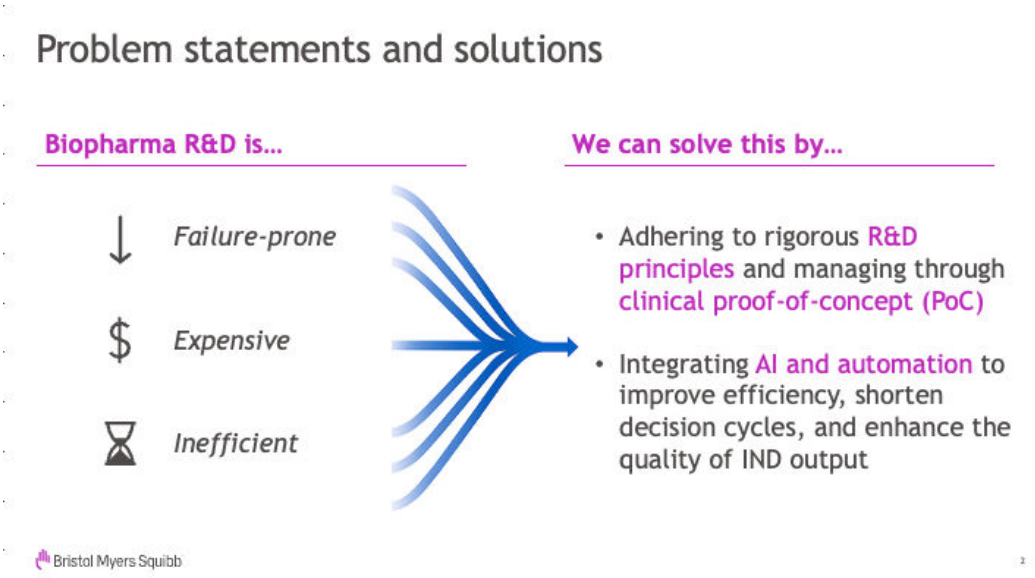


Figure 1. Problem statements and solutions (slide 2).

Context from the literature. Each problem is well documented:

- **Failure-prone:** only about 10–14% of drug programs entering Phase 1 reach approval, and the rate is far lower in some therapeutic areas (Wong, Siah and Lo, Biostatistics 2019). Most late-stage failures are for lack of efficacy, and roughly a quarter for safety (Harrison, Nat Rev Drug Discov 2016).
- **Expensive:** the average cost to take an asset from discovery to launch at the top 20 biopharma companies reached about USD 2.67B in 2025 (Deloitte, 2025).
- **Inefficient:** much of the industry works on the same targets. About 92% of human protein-coding genes have never had a clinical-stage program, while 81 genes (0.4% of the genome) account for 44% of all clinical programs since 2000 (Citeline data released by Minikel et al., Nature 2024; analysis prepared for this symposium). The number of novel targets entering the pipeline each year has fallen by about 70% over the past decade (L.E.K. Consulting, 2025).

2. End-to-end R&D principles

Five principles span the R&D continuum (Figure 2): **causal human biology**, **matching modality to mechanism**, **path to clinical PoC**, accelerated full development, and maximize market access. Together they create a common language to evaluate and prioritize programs and to minimize misalignment between research, development and commercial teams.

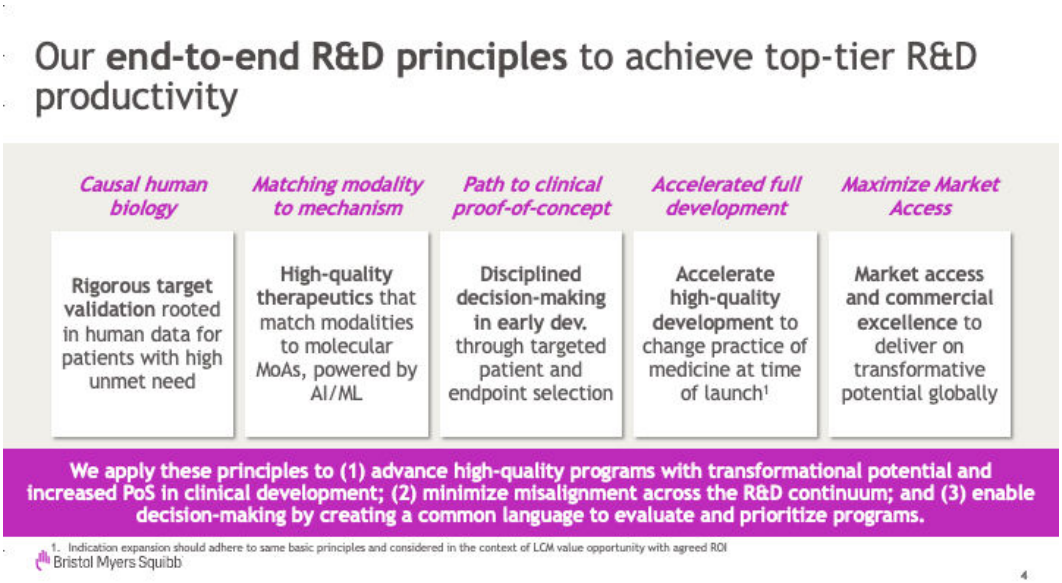


Figure 2. *End-to-end R&D principles (slide 4).*

Human genetics enters the first three principles directly (Table 1).

The evidence that genetics raises success rates has strengthened over a decade. Drug mechanisms with genetic support were about twice as likely to reach approval (Nelson et al., Nat Genet 2015), with stronger effects when the causal gene is confidently assigned (King, Davis and Degner, PLoS Genet 2019). The most recent analysis estimates a 2.6-fold higher probability of success from Phase 1 to launch. The benefit varies by therapeutic area but is largely independent of genetic effect size or allele frequency (Minikel et al., Nature 2024). Principle-based frameworks also measurably improve outcomes. After AstraZeneca adopted its 5R framework (right target, tissue, safety, patient and commercial potential), the share of its projects that completed Phase 3 rose from 4% to 19% (Morgan et al., Nat Rev Drug Discov 2018).

Table 1. *Where human genetics enters the first three R&D principles.*

Principle	What it asks	Contribution of human genetics
Causal human biology	Is the target causal in human disease, in which direction and by how much?	Germline and somatic genetics provide causal evidence; genetic support raises PoS about 2.6-fold
Matching modality to mechanism	Does the modality deliver the pharmacology the genetics implies?	Direction of effect (inhibit, degrade, activate) and the level of modulation needed
Path to clinical PoC	Can a decisive PoC be run quickly and cheaply?	Genetically defined patients, endophenotypes and quantitative biomarkers

Key message

R&D principles create a common language for decisions. Human genetics is the strongest single input to causal human biology and also shapes modality choice and the path to PoC.

3. Where is the most value created? The CVCP

The talk poses a question: what is the inflection point that creates the most value in biopharma? **The answer is the path to clinical PoC**, defined as the critical value creation period (Figure 3).

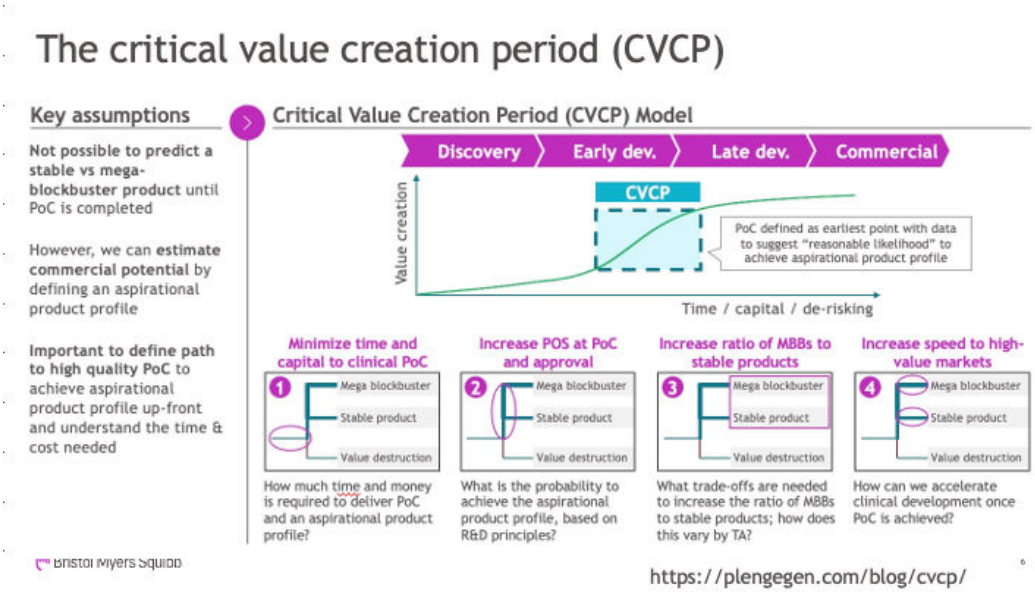


Figure 3. The critical value creation period (slide 6). Source: plengegen.com/blog/cvcp.

Three assumptions underpin the model:

- It is **not possible to predict a stable product versus a mega-blockbuster until PoC is completed.**
- Commercial potential can still be estimated up front by defining an aTPP.
- The path to a high-quality PoC, including its time and cost, should be defined in advance.

PoC is defined as the earliest point at which data suggest a reasonable likelihood of achieving the aspirational product profile. Before it, programs have little value relative to an approved medicine. After it, value rises steeply and the organization can decide how hard to accelerate and how many indications to pursue.

Why PoC is the inflection point. Phase 2 has the lowest transition probability of any development phase. Modeling of R&D productivity has shown that cutting Phase 2 attrition, by stopping weak programs early and cheaply, improves the cost per launch more than almost any other lever (Paul et al., Nat Rev Drug Discov 2010). The CVCP formalizes that insight: invest to reach a decisive PoC quickly, then concentrate capital on programs that pass.

3.1 Why mega-blockbusters matter

This section gives context that the talk does not cover in detail. It draws on external sources rather than the slides.

In this guide, **mega-blockbusters** are products with **more than USD 5B in annual revenue**, and **stable products** are those with **USD 3B or less**. Many stable products are important medicines for patients; the distinction here is financial. Three features of biopharma economics make mega-blockbusters essential for a large company.

(a) Returns are extremely skewed. Studies of US new drug introductions have repeatedly found that a small fraction of drugs generates most of the value. In the 1988 to 1992 cohort, the top decile of new drugs accounted for 56% of total sales revenue, and peak sales of the top decile were several times those of the next decile (Grabowski and Vernon, 2000). Earlier work found that only the top 30 of 100 drugs introduced in the 1970s covered average

R&D costs on a fully allocated basis (Grabowski and Vernon, 1990). The pattern holds today. In the 2025 late-stage pipelines of the top 20 biopharma companies, 54 blockbuster assets, only 9% of drugs in advanced development, are forecast to deliver about 70% of projected sales. About 53% of pipeline asset-indications have forecast peak sales below USD 250M (Deloitte, 2025).

(b) Average R&D returns barely clear the cost of capital, and depend on a few products. Deloitte's projected internal rate of return (IRR) on late-stage pipelines rose to 7.0% in 2025. Excluding one drug class (GLP-1 and GIP agonists for obesity and diabetes), it falls to 2.9%, and the median company IRR was about 4.0%. Average forecast peak sales per asset were USD 598M, or USD 353M without that class. The average cost to take an asset from discovery to launch reached about USD 2.67B (Deloitte, 2025). **Without mega-blockbusters, the average program does not earn back its cost of capital.**

(c) Revenue must be replaced before it can grow. Large companies lose most of a product's revenue within a few years of losing exclusivity. Between 2025 and 2030, estimates of branded revenue facing loss of exclusivity range from about USD 180B to more than USD 300B, roughly one-sixth of industry sales, including about 70 blockbusters (Evaluate; EY; GlobalData; PitchBook, as summarized in industry reviews). A large company typically gains only about one to two new approvals per year: a top 20 company averaged about one new FDA-approved drug per year from 2014 to 2023, including acquisitions, and the most productive reached about 2.3 (Schuhmacher et al., 2024). With so few launches, each must be large.

Illustrative growth arithmetic. Consider a hypothetical company with USD 50B in annual revenue that aims to grow over 10 years while part of its base loses exclusivity (Table 2). Assume stable products average USD 1B in peak sales (generous, given the USD 598M industry average) and mega-blockbusters average USD 7B. The share of launches that must become mega-blockbusters is then:

Table 2. *Illustrative share of new launches that must reach mega-blockbuster status, for a hypothetical company with USD 50B revenue over 10 years.*

Revenue growth (CAGR)	Base revenue lost to loss of exclusivity	New revenue needed by year 10 (USD B)	Share of 25 launches that must be mega-blockbusters	Share of 15 launches that must be mega-blockbusters
3%	30%	32	5%	19%
3%	50%	42	11%	30%
5%	30%	46	14%	35%
5%	50%	56	21%	46%

New revenue needed = 50 x (1 + CAGR)^10 minus the retained base. Required share f solves launches x [1 + 6f] = new revenue, with stable products at USD 1B and mega-blockbusters at USD 7B. Simplified: assumes all launches reach peak sales by year 10 and ignores pricing changes and acquisitions of marketed products.

For a company launching about 2 to 2.5 new medicines a year, roughly 10–20% of launches must become mega-blockbusters to sustain mid-single-digit growth through a typical loss-of-exclusivity cycle. At about 1.5 launches a year, the required share rises to 20–45%, which is unrealistic. This is why the ratio of mega-blockbusters to stable products is one of the four CVCP actions, and why the CVCP matters. Mega-blockbuster potential cannot be identified until PoC, so a portfolio must carry enough programs through PoC, quickly and cheaply, to find them.

3.2 Four actions

Because mega-blockbusters and stable products cannot be told apart until PoC, four actions follow (Table 3).

Table 3. The four CVCP actions and how human genetics contributes.

Action	Question to answer	Contribution of human genetics
1. Minimize time and capital to PoC	How much time and money does it take to deliver PoC and the aTPP?	Genetic patient selection and biomarkers allow smaller, faster PoC studies
2. Increase PoS at PoC and approval	What is the probability of achieving the aTPP, based on R&D principles?	Causal genetic evidence raises the starting PoS
3. Increase the ratio of mega-blockbusters to stable products	What trade-offs are needed, and how do they vary by therapeutic area?	Genetic breadth (the same pathway across many diseases) signals mega-blockbuster potential
4. Increase speed to high-value markets	How can development accelerate once PoC is achieved?	Shared genetic architecture ranks concurrent indications

Key message

The most value is created on the path to clinical PoC. A portfolio that needs mega-blockbusters must reach PoC quickly, cheaply and decisively, because only PoC reveals which programs have that potential.

4. Four examples of programs that navigate the CVCP

Programs differ in the time and investment needed to reach PoC and in their commercial potential (Figure 4). The approach must be tailored to each:

- **(a) Early evidence of mega-blockbuster potential:** invest capital to accelerate.
- **(b) Long time or high investment to determine mega-blockbuster potential:** optimize the path to the aTPP.
- **(c) Stable product potential:** reach PoC cost-efficiently.
- **(d) Long time or high investment for a stable product:** consider termination.

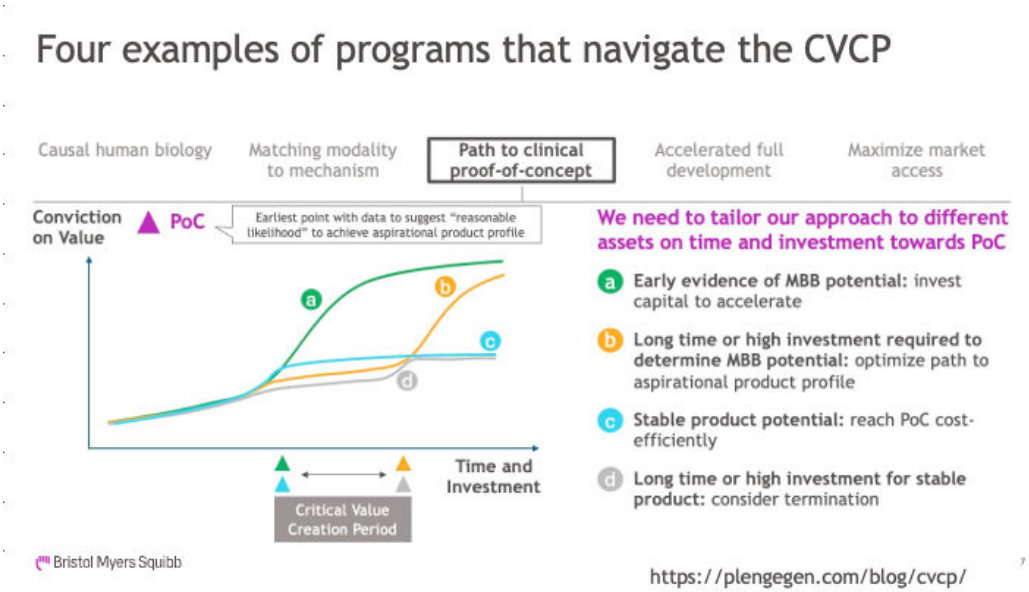


Figure 4. Four hypothetical programs navigating the CVCP (slide 7). Source: plengegen.com/blog/cvcp.

Genetics is one of the main tools for moving a program from (b) toward (a), or from (d) toward (c). It does so mainly by defining patient subgroups in which a small, fast PoC is unambiguous.

5. Monitoring CVCP inflection points

5.1 New axes: PoS against time and cost to PoC

To manage individual programs, the CVCP is redrawn with **time and investment on the X-axis** and **probability of success on the Y-axis** (Figure 5). A typical program moves through a series of inflection points: development candidate (DC), toxicology results, IND, safety and tolerability, PK-PD, early clinical data and PoC. Each raises PoS by a different amount at a different cost. The steepest gains come late, at early clinical data and PoC.

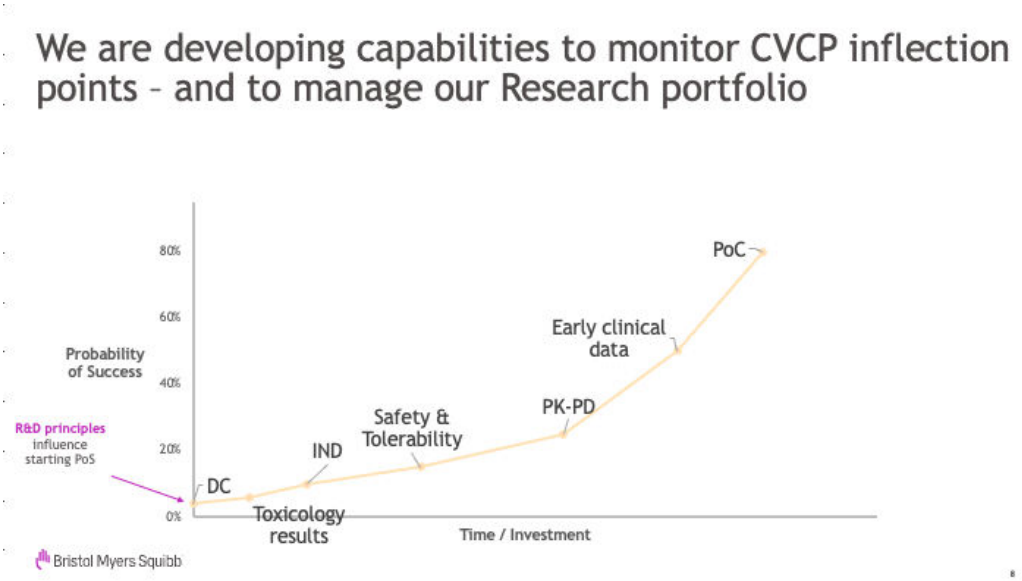


Figure 5. Inflection points on the path to PoC (slide 8). R&D principles set the starting PoS.

5.2 Every program is different

Programs differ in three ways (Figure 6): **starting PoS**, **the inflection points they pass through**, and **the time and capital needed to reach PoC**. A fast path to PoC is ideal regardless of commercial potential. A medium path is acceptable for medium and large commercial opportunities. A slow path is acceptable only for large ones. Competitor data can also move a program's PoS, alongside its own internal data.

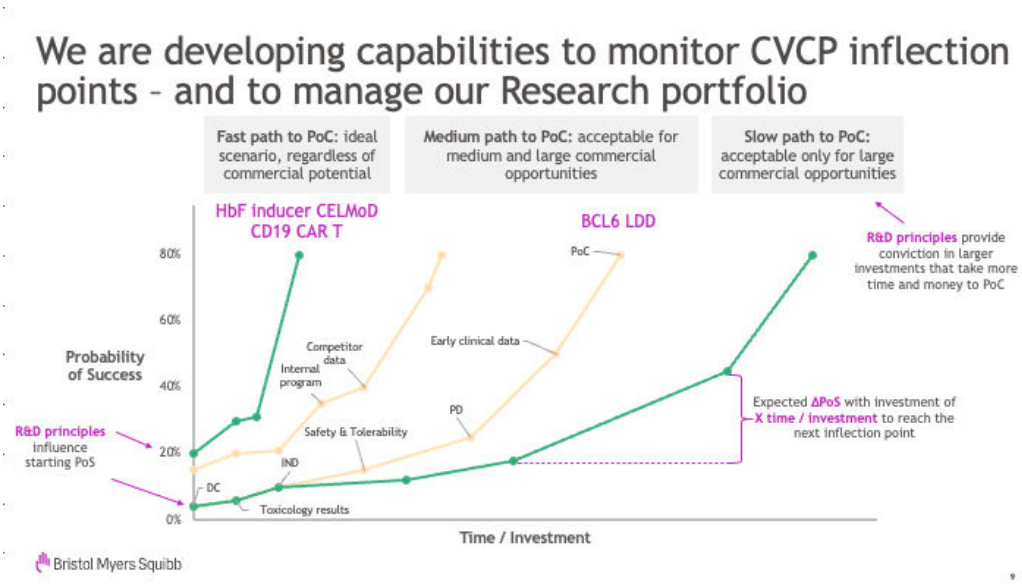


Figure 6. Different paths to PoC, with the three example programs placed on fast and medium paths (slide 9).

Three key messages

- **R&D principles influence the starting point.** Strong causal human biology, often from genetics, starts a program higher on the PoS axis and gives conviction for larger investments that take more time and money to reach PoC.
- **Dollars spent per change in PoS is the measure that matters.** Each inflection point should be judged by the expected change in PoS for the time and investment needed to reach it.
- **Revenue potential, set by the aTPP, is refined through PoC.** The commercial case is a hypothesis until PoC data show whether the aspirational profile is achievable.

5.3 How genetics changes PoS before and after PoC

Genetics acts on the curve in two ways. It **raises the starting PoS**: drug mechanisms with human genetic support are about 2.6 times more likely to succeed (Minikel et al., Nature 2024). It also **makes the PoC inflection point more decisive** by supplying biomarkers, dose estimates and patient selection. Table 4 and Figure 7 illustrate the combined effect with assumed values.

Table 4. Illustrative probability of approval before and after PoC.

Scenario	Before PoC	After positive PoC	After negative PoC
No genetic support, weak PoC	0.10	0.16	0.06
No genetic support, decisive PoC	0.10	0.41	0.02
Genetic support, weak PoC	0.26	0.37	0.17
Genetic support, decisive PoC	**0.26**	**0.58**	**0.06**

Illustrative model: P(approval given a true hypothesis) = 0.7; genetic support about 2.6-fold; decisive PoC sensitivity 0.85 and specificity 0.90; weak PoC 0.65 and 0.65. Values support relative comparison only.

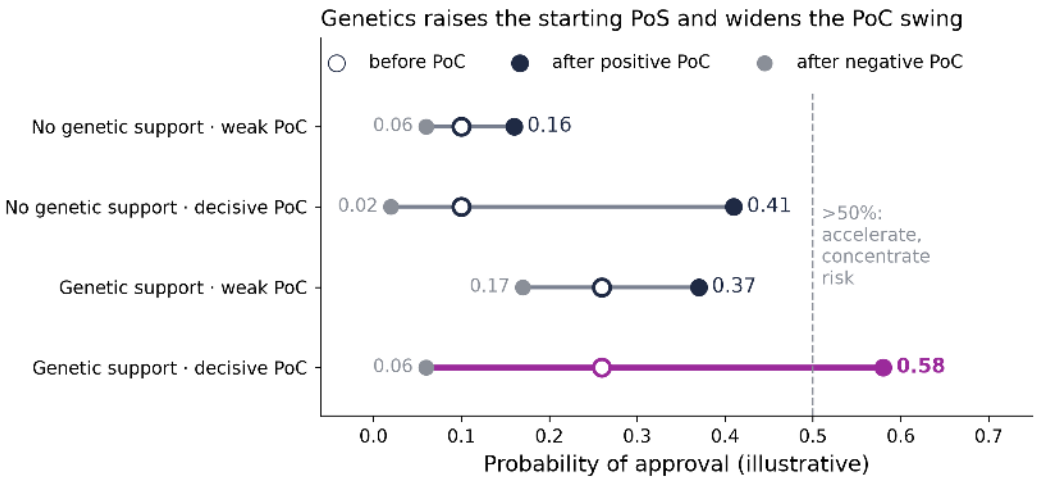


Figure 7. *Genetics raises the starting PoS and widens the swing at PoC. Open circles: before PoC; filled circles: after a positive PoC (dark) or a negative PoC (grey). Same values as Table 4.*

Only genetic support combined with a decisive PoC reaches the post-PoC range above 50% that justifies acceleration. A genetically supported program with a weak PoC is the costliest case: capital is spent and no clear decision results.

6. Three programs on the path to PoC

Each example is read through the three R&D principles and its position on the CVCP curve (Figure 6). The programs differ in **where their causal evidence comes from**: germline genetics, somatic cancer genetics, or human pharmacology (Table 5). That difference shapes each path to PoC.

6.1 HbF-inducing CELMoD for sickle cell disease (BMS-986470)

The program. BMS-986470 is an oral molecular glue degrader (CELMoD) that selectively degrades ZBTB7A and WIZ, two repressors of fetal hemoglobin (HbF). It recapitulates human genetics and has shown preclinical benefit ex vivo in blood from patients with sickle cell disease (SCD) (Figure 8). A Phase 1/2a PoC study is enrolling (NCT06481306).

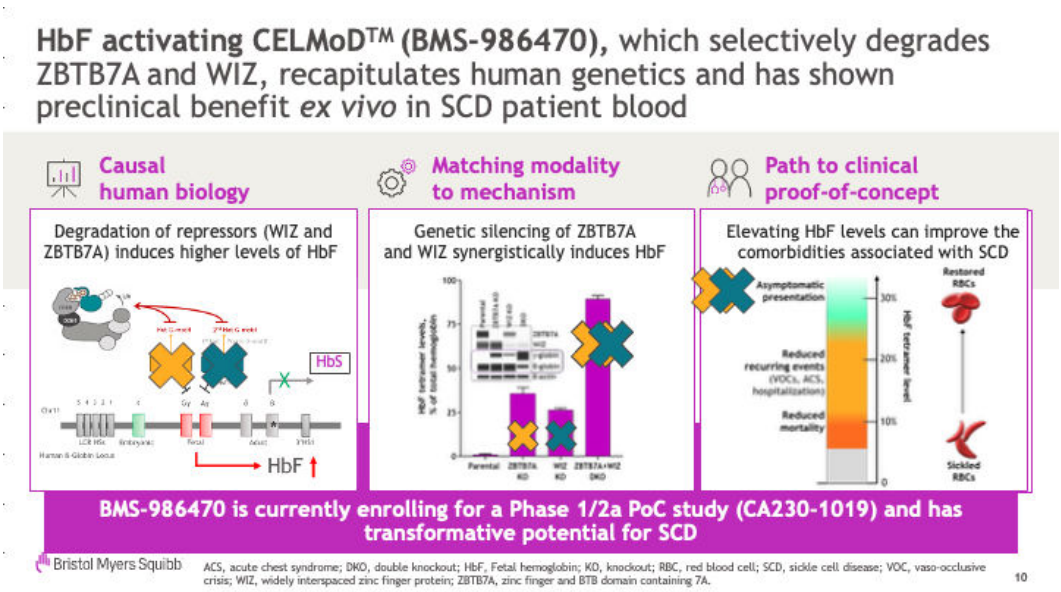


Figure 8. BMS-986470 on the three R&D principles (slide 10). Higher HbF levels are associated with reduced mortality, fewer recurring events and, at the highest levels, asymptomatic presentation.

Scientific background: germline genetics points to HbF. SCD is caused by a single amino-acid change in adult beta-globin. Fetal hemoglobin does not take part in sickle polymerization, so raising HbF dilutes the pathogenic hemoglobin.

- **Clinical observation:** higher HbF levels are associated with longer survival in SCD (Platt et al., NEJM 1994).
- **Rare variants:** people who inherit hereditary persistence of fetal hemoglobin (HPFH) alongside the sickle mutation often have mild disease. This is a natural experiment showing that lifelong high HbF is protective and well tolerated.
- **Common variants:** genome-wide association studies identified common variants at BCL11A and other loci that explain much of the population variation in HbF (Uda et al., PNAS 2008). Fine-mapping located the key variants in an erythroid-specific BCL11A enhancer (Bauer et al., Science 2013).
- **From genetics to medicine:** editing that enhancer in patients' own stem cells raised HbF and eliminated vaso-occlusive crises in early clinical studies (Frangoul et al., NEJM 2021). That work led to an approved gene-editing therapy, clinical proof that genetically guided HbF induction works.

Scientific background: the two targets have different sources of validation.

- **ZBTB7A (also called LRF) has germline support.** It was identified as a direct repressor of fetal globin, independent of BCL11A (Masuda et al., Science 2016). Naturally occurring HPFH point mutations about 200 base pairs upstream of the fetal globin genes disrupt ZBTB7A binding and raise fetal globin expression (Martyn et al., Nat Genet 2018). These rare variants are, in effect, human genetic proof that losing ZBTB7A repression at this site raises HbF.
- **WIZ was found by chemical, not human, genetics.** A phenotypic screen of cereblon-binding molecular glues in human erythroblasts revealed WIZ as a previously unknown HbF repressor, and degrading WIZ raised HbF in humanized mice and monkeys (Ting et al., Science 2024). BMS-986470 also emerged from a phenotypic screen in primary erythroblasts. Dual silencing of both repressors raised the proportion of HbF-expressing red cells above 95% and total HbF above 80% (Blood 2024; 144(Suppl 1): 169). A recent review places these degraders within the wider HbF field (Biomarker Research 2026).

CVCP read: HbF CELMoD (fast path)

- **Causal human biology:** germline genetics (HPFH and common HbF variants) validates raising HbF, and rare HPFH mutations validate ZBTB7A specifically; WIZ adds synergy found through chemical genetics.
- **Path to PoC:** genetics also sets the target level. HbF is a quantitative biomarker with known clinical thresholds (about 30% HbF distributed across most red cells is predicted to prevent sickling), so PoC can be read early. This places the program on a fast path.
- **aTPP:** an oral therapy with transformative potential for SCD, including in regions where cell-based therapies are impractical.

6.2 BCL6 ligand-directed degrader for B-cell lymphoma (BMS-986458)

The program. BMS-986458 is an oral heterobifunctional ligand-directed degrader (LDD) that links a BCL6-binding ligand to cereblon (Figure 9). It is highly selective across the human proteome: BCL6 is the only protein significantly degraded. Preclinically, exposure in plasma and tumor drives near-complete BCL6 degradation (J Med Chem 2026). About 30% of diffuse large B-cell lymphomas (DLBCL) show high BCL6 expression by immunohistochemistry (IHC). A Phase 1/2 study is under way in relapsed or refractory non-Hodgkin lymphoma (NCT06090539). Early dose-escalation results showed rapid BCL6 degradation in blood and tumor and objective responses in heavily pretreated patients with DLBCL and follicular lymphoma, with updates at ASH 2025 (BMS, December 2025).

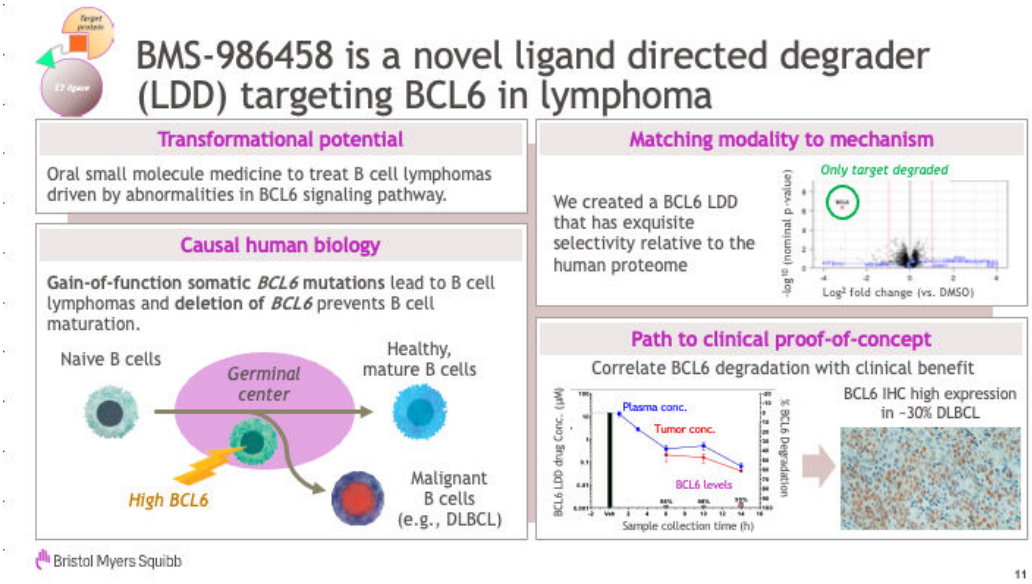


Figure 9. BMS-986458 on the three R&D principles (slide 11).

Scientific background: somatic genetics points to BCL6. BCL6 is a transcriptional repressor that is essential for germinal center B cells. Mice lacking BCL6 cannot form germinal centers (Ye et al., Nat Genet 1997; Dent et al., Science 1997).

- **Somatic genetics:** in lymphoma, BCL6 is deregulated by chromosomal translocations and promoter mutations. BCL6 alterations help define genetic subtypes of DLBCL (Schmitz et al., NEJM 2018; Chapuy et al., Nat Med 2018).
- **Experimental proof of causality:** mice engineered to express BCL6 constitutively develop a disease resembling human DLBCL (Cattoretti et al., Cancer Cell 2005).

- **Druggability:** BCL6 has long been considered hard to drug. Small molecules that bind its BTB domain can trigger its degradation (Slabicki et al., Nature 2020), which supports degradation as the matching modality.

This is a case where **somatic rather than germline genetics** supplies the causal evidence, as is typical in oncology. The genetic lesion is in the tumor, so it also defines the patients most likely to respond.

CVCP read: BCL6 LDD (medium path)

- **Causal human biology:** somatic cancer genetics (translocations and mutations that deregulate BCL6) and engineered mouse models support causality.
- **Matching modality to mechanism:** a degrader that removes only BCL6 matches the mechanism the genetics implies and addresses a target that was hard to drug with inhibitors.
- **Path to PoC:** tumor degradation is a direct measure of target engagement, and IHC selects the roughly 30% of patients most likely to benefit. The next inflection point is to correlate degradation with durable clinical benefit.

6.3 CD19 cell therapy for immune reset in autoimmune disease

The program. CD19 CAR T approaches aim to induce immune reset across multiple autoimmune diseases (Figure 10). The program spans autologous (zola-cel), allogeneic (BMS-986515) and in vivo (BMS-986553) CD19 CAR T. Zola-cel produced robust B-cell depletion with recovery of naive B cells across autoimmune indications. Pivotal trials of zola-cel are under way in SLE and systemic sclerosis, a healthy-donor allogeneic Phase 1 study is enrolling, and an in vivo CD19 CAR T IND is planned for 2026.

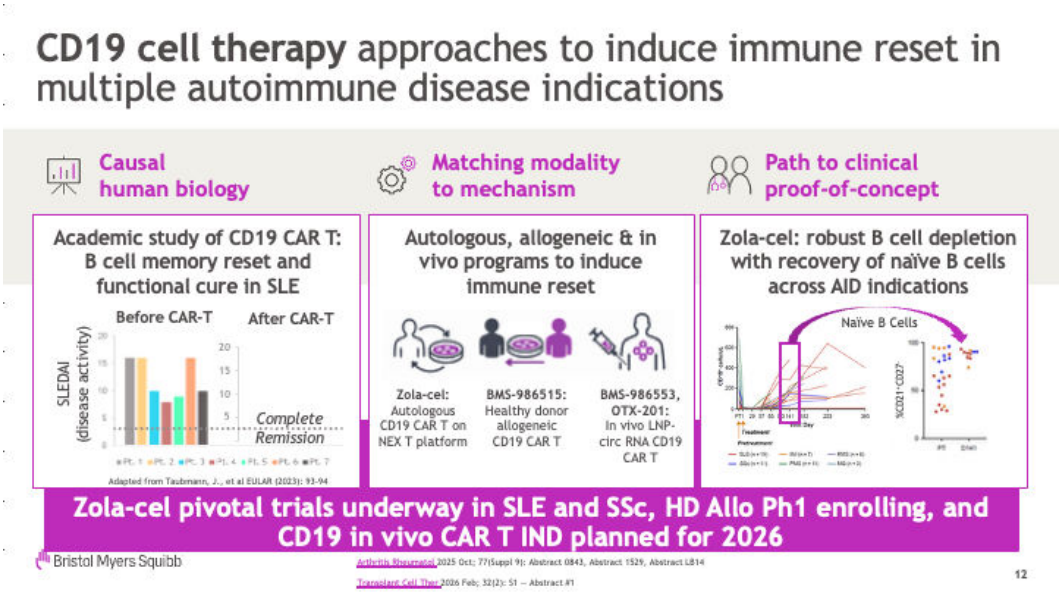


Figure 10. CD19 cell therapy on the three R&D principles (slide 12).

Scientific background: human pharmacology leads, genetics supports. CD19 is a B-cell co-receptor present from early B-cell development through plasmablasts.

- **Human pharmacology:** the decisive human evidence came from treating patients. In refractory systemic lupus erythematosus (SLE), CD19 CAR T cells produced drug-free remission (Mackensen et al., Nat Med 2022). In a case series of 15 patients with SLE, myositis or systemic sclerosis, all patients stopped immunosuppression. B cells returned after a mean of about 112 days, predominantly as naive cells, without relapse during follow-up

(Müller et al., NEJM 2024). Antibody-based B-cell depletion has given mixed results in SLE, which suggests that the depth of tissue depletion matters.

- **Germline genetics as a human knockout:** people with inherited CD19 deficiency have an antibody-deficiency syndrome but otherwise develop normally (van Zelm et al., NEJM 2006). This defines the expected on-target consequence of prolonged CD19 loss.
- **Polygenic genetics:** genome-wide association studies of SLE implicate many genes in B-cell signaling and activation (Bentham et al., Nat Genet 2015), supporting a causal role for B cells across the disease.

CVCP read: CD19 CAR T (fast path)

- **Causal human biology:** human pharmacology, not genetics, supplied the first causal evidence that resetting B cells can induce remission. Germline genetics adds a human-knockout view of on-target safety and polygenic support for B-cell involvement.
- **Path to PoC:** immune reset is visible in a small number of patients, which gives early evidence of mega-blockbuster potential and justifies investing to accelerate.
- **aTPP:** large, because the mechanism may apply across many autoimmune diseases. The genetics opportunity is to rank concurrent indications by shared B-cell genetic architecture and to identify patients likely to achieve durable reset.

Table 5. Sources of human evidence for each target.

Target	Germline genetics	Somatic genetics	Other human or experimental evidence	What the evidence informs
HbF axis (BMS-986470)	Strong: HPFH; common variants at BCL11A and other loci	Not applicable	Survival tracks HbF level; editing the BCL11A enhancer is clinically effective	Target direction, target HbF level, population
ZBTB7A	Supportive: rare HPFH promoter mutations disrupt ZBTB7A binding	Not applicable	Knockout raises HbF	Mechanism confirmed by natural variation
WIZ	None reported	Not applicable	Chemical genetics: phenotypic degrader screen; knockout	Synergy with ZBTB7A degradation
BCL6 (BMS-986458)	Not a germline target	Strong: translocations, promoter mutations, genetic DLBCL subtypes	Engineered mouse models; high BCL6 by IHC in about 30% of DLBCL; early clinical activity	Patient selection, pharmacodynamic marker
CD19 (CAR T)	Human knockout (CD19 deficiency); polygenic SLE risk in B-cell genes	Not applicable	Human pharmacology: remission after CD19 CAR T	Safety, indication breadth

Table 6. The three programs compared.

Program	Source of causal evidence	Path to PoC	Decisive readout	Genetic lever
HbF CELMoD (SCD)	Germline HbF genetics	Fast	HbF level against known clinical thresholds	Target, direction and dose
BCL6 LDD (DLBCL)	Somatic BCL6 genetics	Medium	Tumor degradation linked to clinical benefit	Patient selection by BCL6 level
CD19 CAR T (autoimmune)	Human pharmacology	Fast	Immune reset in few patients	Indication ranking and segmentation

Key message

Three programs, three kinds of causal evidence. **Germline genetics** fixed the target and the biomarker for the HbF program; **somatic genetics** defines both causality and the patients for the BCL6 program; **human pharmacology** established CD19 immune reset, with germline genetics adding safety and indication context. Knowing the source of evidence tells a portfolio team which inflection point carries the most risk.

7. Portfolio view

The same dimensions can be applied across a whole portfolio (Figure 11). Each program is plotted by **scientific conviction** (the first three R&D principles) against **CVCP value** (change in PoS from first-in-human to PoC per dollar invested). Bubble size shows the aTPP. Each quadrant implies a default action:

- **Accelerate:** fund fully and plan concurrent indications after PoC.
- **Optimize the path:** reduce the cost and time to PoC with genetic segmentation and biomarkers.
- **Build conviction:** run killer experiments before spending on PoC.
- **Stop or redesign:** redeploy capital to the upper right.

Why the portfolio view matters. Value is concentrated: about 9% of late-stage assets are forecast to deliver about 70% of projected sales (Section 3.1). Targets are also crowded. When many programs pursue the same target, one readout can move all of them at once, so crowded portfolios carry correlated risk. Genetically supported novel targets offer both higher starting PoS and diversification.

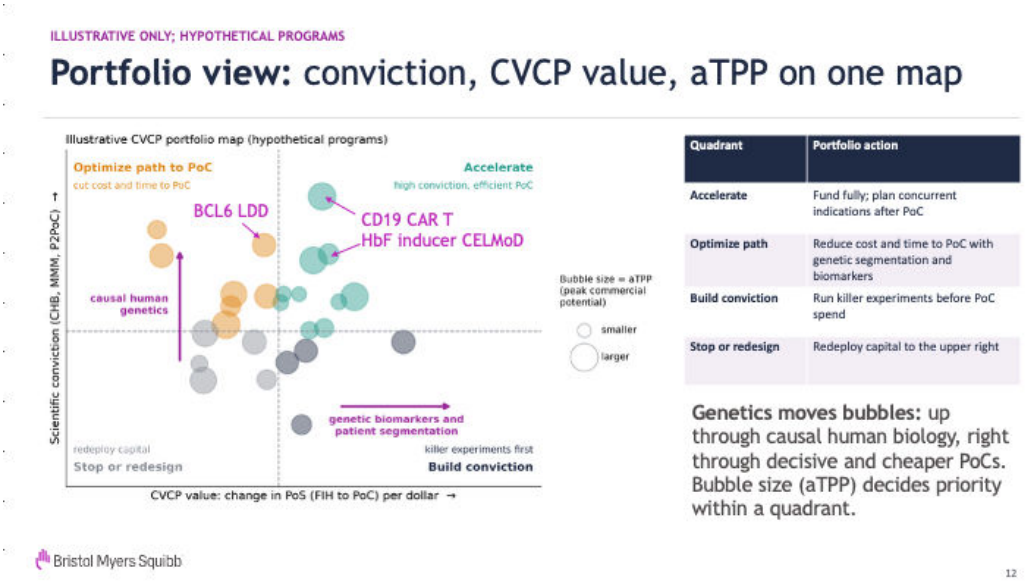


Figure 11. Illustrative portfolio map with hypothetical programs and the three example programs indicated (slide 13).

> **Key message**
>
> **Genetics moves bubbles:** up through causal human biology, and right through more decisive and cheaper PoCs. Within a quadrant, bubble size (aTPP) sets priority. Crowded targets are correlated bets, because one readout can move many bubbles at once.

8. AI and automation: dream and build

The second solution from Section 1 is to integrate AI and automation. BMS is **investing in the dream state** and **building the individual components** needed to realize it.

Industry context. Investment in AI drug discovery is rising quickly but is concentrated in molecule design, where data are structured and feedback is fast. Of the 15 largest pharma–AI discovery deals from 2023 to mid-2026, 14 included molecule design, 4 included target or disease biology, and none addressed preclinical-to-clinical translation. AI for target identification is rated only of medium maturity, partly because models struggle to separate causation from correlation (McKinsey & Company, *AI drug discovery: focusing on what matters most*, September 2026). The CVCP framework points AI at the opposite end of the value chain: causal biology and the path to PoC.

8.1 The dream: an AI-enabled research organization

The aspiration spans the first three R&D principles (Figure 12):

- **Causal human biology:** complete prediction of the safety and efficacy profile of every target nominated in early discovery, from a causal foundation model that integrates multimodal data, in silico experiments and AI inference.
- **Matching modality to mechanism:** development candidates generated through a fully automated, closed-loop laboratory.
- **Path to clinical PoC:** dynamic management of every clinical experiment through the CVCP and predictive modeling.

Reverse translation connects the steps: AI analyzes clinical feedback and external data to update mechanistic understanding and improve the models. The intended results are an expanded portfolio with better prioritization, faster cycles with fewer iterations, and higher PoS in development.

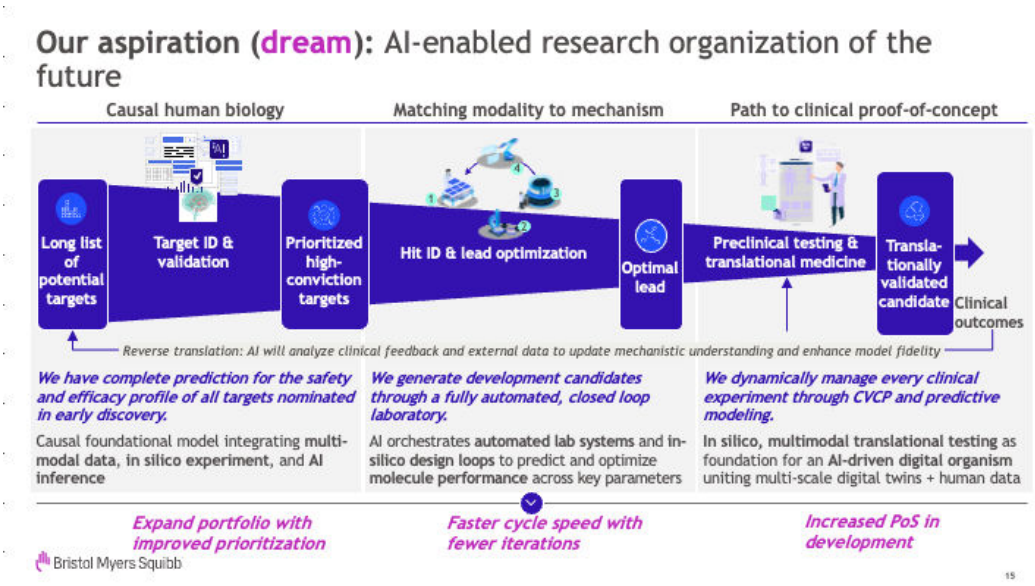


Figure 12. *The aspiration: an AI-enabled research organization of the future (slide 15).*

8.2 The build: five interdependent capabilities

Five capabilities are being built to realize that vision (Figure 13, Table 7).

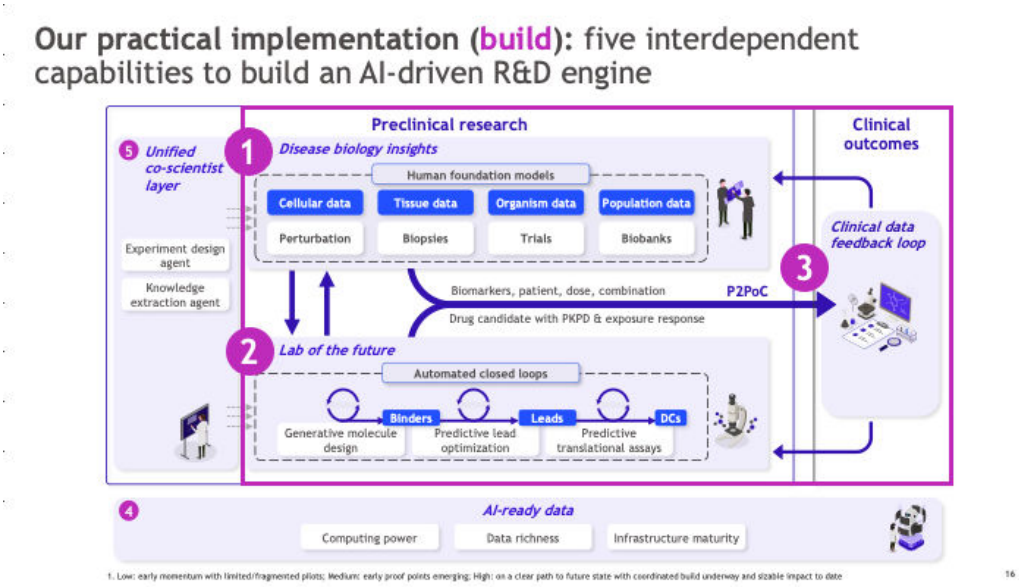


Figure 13. Practical implementation: five interdependent capabilities for an AI-driven R&D engine (slide 16).

Table 7. The five capabilities and their role in the CVCP.

Capability	What it includes	Role in the CVCP
1. Disease biology insights	Human foundation models built on cellular (perturbation), tissue (biopsy), organism (trial) and population (biobank) data	Raises starting PoS; population genetics is a direct input
2. Lab of the future	Automated closed loops for generative molecule design, predictive lead optimization and predictive translational assays	Cuts time and capital to a development candidate
3. Clinical data feedback loop	Biomarkers, patients, dose and combinations for the path to PoC; drug candidates with PK-PD and exposure response	Makes PoC decisive and feeds results back into models
4. AI-ready data	Computing power, data richness, infrastructure maturity	Foundation for all capabilities
5. Unified co-scientist layer	Agents for experiment design and knowledge extraction	Keeps CVCP assessments current across programs

Key message

AI scales the R&D principles and the CVCP. Human genetics, in the form of population biobank data and causal foundation models, is central to the first capability and to the feedback loop that recalibrates PoS after each readout.

8.3 An external example of the future state: DrugTargetWorld

A recent preprint from Stanford University, **DrugTargetWorld**, shows what it could look like to train and test AI agents on the causal-biology part of the CVCP (Margolis et al., arXiv 2026). It addresses a basic obstacle: real biobanks cannot tell an AI agent whether it found the right target, because true causal drivers are not known in advance and confirming them takes experiments.

How it works.

- **Synthetic biobanks with a hidden answer.** The framework generates simulated biobanks, or "worlds", from a concealed causal model. Each world contains genotypes, proteins, health records and outcomes for 54,000

participants, including 10,800 with synthetic cardiac MRI. The causal model specifies which proteins truly drive disease and in which direction. The current version covers cardiovascular disease (dilated, hypertrophic and ischemic heart disease, and heart failure with preserved ejection fraction), and the framework is designed to extend to other diseases.

- **Planted biases.** Each world can include up to nine recognized ways that observational data mislead, including confounding, reverse causation, selection bias, genetic pleiotropy, imaging batch effects, and a "harmful surrogate". The harmful surrogate is a protein that improves an imaging measure without being causal, and advancing it is penalized.
- **An end-to-end task.** Agents must build a disease phenotype from raw data, nominate causal proteins, state whether each should be inhibited or activated, and explain why misleading proteins should be rejected.
- **Budgeted experiments.** Agents can buy virtual experiments (cell perturbations or longer-term knockdowns) under budgets of zero, USD 450,000 or USD 2 million. This forces them to judge when an experiment is worth its cost.
- **Scoring against ground truth.** Scoring is out of 100 points: 30 for target identification, 25 for confidence matched to evidence, 20 for direction of effect, 15 for identifying biases and 10 for the disease phenotype. The authors describe this as a verifiable reward, which also allows agents to be trained, not just evaluated.

What the authors found.

- **Analyses were within reach.** Across 540 episodes with nine agents, the two leading agents recovered about 64% of causal drivers on average. They based most nominations on cis Mendelian randomization and usually checked instrument strength.
- **Integrative judgments were not.** No agent reliably rejected misleading non-causal proteins; the best averaged under 3 of 15 points for bias identification. Leading agents scored about 35 to 40 of 100 overall.
- **Experiments were used selectively.** A larger experimental budget did not detectably change the leading agents' scores. When two proteins shared a genetic instrument, the leading agents often ran an experiment on the causal one, but no agent tested a protein whose effect appeared only over long-term follow-up.
- **The authors' conclusion:** current agents can perform most of the individual analyses in a biobank study but do not yet consistently make the integrative judgments that connect them. These are how to measure disease, how to separate causal drivers from non-causal proteins, and when evidence is sufficient to support a claim.

Why it is relevant to the CVCP. The benchmark reproduces, in miniature, the decisions that set a program's starting PoS and path to PoC (Table 8). Its emphasis on direction of effect, calibrated confidence and the cost of experiments mirrors the CVCP measure of dollars spent per change in PoS. Environments like this could help build and test the "causal foundation model" and co-scientist agents described in Section 8.1, before they are trusted with real portfolio decisions.

Table 8. *How the DrugTargetWorld task maps onto the R&D principles and the CVCP.*

DrugTargetWorld element	R&D principle or CVCP concept	Real-world analog
Identify causal driver proteins	Causal human biology	Target validation with human genetics (for example, Mendelian randomization)
Infer direction of modulation	Matching modality to mechanism	Choosing to inhibit, degrade or activate
Construct a disease phenotype from raw data	Path to clinical PoC	Choosing endpoints and endophenotypes
Reject misleading proteins with the right bias	Causal human biology	Avoiding confounded or pleiotropic targets
Penalty for advancing a harmful surrogate	Path to clinical PoC	Surrogate endpoints that do not predict outcomes
Buy experiments under a budget	CVCP value: dollars per change in PoS	Deciding which killer experiment to fund
Confidence matched to evidence	Portfolio decision-making	Go/no-go governance

9. Conclusions

Conclusions

- **Problem statement:** drug R&D remains failure-prone, expensive and inefficient.
- **R&D principles and the CVCP:** high-conviction targets paired with the right modality unlock novel biology and clinical differentiation. The most value is created on the path to clinical PoC, and programs should be managed by dollars spent per change in PoS.
- **AI and automation:** scale these principles to improve speed, quality and PoS across the pipeline, through both the dream and the build.

What this asks of human geneticists. The CVCP benefits most from genetic outputs that change the shape of a program's path. These include direction and size of effect with uncertainty; allelic series that indicate the level of modulation needed; quantitative, genetically anchored biomarkers that read out early; genetic patient subgroups and endophenotypes; phenome-wide maps that rank concurrent indications; and genotyped trials that feed results back into genetic priors.

References and further reading

Links open the publisher or registry page. References are grouped by the section in which they are discussed.

Notes on sources. Figures 1 to 6 and 8 to 13 are slides from *Managing an R&D portfolio through the critical value creation period (CVCP)*, ASHG 2026. The CVCP framework is described at plengegen.com/blog/cvcp (27 March 2025). Table 2, Table 4, Figure 7 and the portfolio map in Figure 11 are illustrative; the map's background programs are hypothetical, and Table 2 is not a forecast for any company. Revenue thresholds for mega-blockbuster and stable products are as defined for this talk. Loss-of-exclusivity estimates (Evaluate, EY, GlobalData, PitchBook) are taken from industry summaries, and the precise figure depends on market scope and dates.

R.1 R&D; productivity and the value of genetics (Sections 1 to 3)

- **[1]** Wong CH, Siah KW, Lo AW. Estimation of clinical trial success rates and related parameters. *Biostatistics* 2019. doi:10.1093/biostatistics/kxx069
- **[2]** Harrison RK. Phase II and phase III failures: 2013-2015. *Nat Rev Drug Discov* 2016. doi:10.1038/nrd.2016.184
- **[3]** Paul SM, et al. How to improve R&D productivity: the pharmaceutical industry's grand challenge. *Nat Rev Drug Discov* 2010. doi:10.1038/nrd3078
- **[4]** Nelson MR, et al. The support of human genetic evidence for approved drug indications. *Nat Genet* 2015. doi:10.1038/ng.3314
- **[5]** King EA, Davis JW, Degner JF. Are drug targets with genetic support twice as likely to be approved? *PLoS Genet* 2019. doi:10.1371/journal.pgen.1008489
- **[6]** Minikel EV, et al. Refining the impact of genetic evidence on clinical success. *Nature* 2024. doi:10.1038/s41586-024-07316-0
- **[7]** Morgan P, et al. Impact of a five-dimensional framework on R&D productivity at AstraZeneca. *Nat Rev Drug Discov* 2018. doi:10.1038/nrd.2017.244
- **[8]** Grabowski HG, Vernon J. The distribution of sales revenues from pharmaceutical innovation. *PharmacoEconomics* 2000. doi:10.2165/00019053-200018001-00005
- **[9]** Grabowski H, Vernon J. A new look at the returns and risks to pharmaceutical R&D. *Management Science* 1990; 36(7).
- **[10]** Deloitte Centre for Health Solutions. *Measuring the return from pharmaceutical innovation: navigating the GLP-1 boom* (2025). deloitte.com
- **[11]** Schuhmacher A, et al. Comparative analysis of FDA approvals by top 20 pharma companies (2014-2023), 2024. researchgate.net
- **[12]** L.E.K. Consulting. *Is biopharma doing enough to advance novel targets?* (2025, updated 2026).

R.2 HbF-inducing CELMoD (Section 6.1)

- **[13]** Platt OS, et al. Mortality in sickle cell disease: life expectancy and risk factors for early death. *NEJM* 1994. doi:10.1056/NEJM199406093302303
- **[14]** Uda M, et al. Genome-wide association study shows BCL11A associated with persistent fetal hemoglobin. *PNAS* 2008. doi:10.1073/pnas.0711566105
- **[15]** Bauer DE, et al. An erythroid enhancer of BCL11A subject to genetic variation determines fetal hemoglobin level. *Science* 2013. doi:10.1126/science.1242088

- **[16]** Frangoul H, et al. CRISPR-Cas9 gene editing for sickle cell disease and beta-thalassemia. *NEJM* 2021. doi:10.1056/NEJMoa2031054
- **[17]** Masuda T, et al. Transcription factors LRF and BCL11A independently repress expression of fetal hemoglobin. *Science* 2016. doi:10.1126/science.aad3312
- **[18]** Martyn GE, et al. Natural regulatory mutations elevate the fetal globin gene via disruption of BCL11A or ZBTB7A binding. *Nat Genet* 2018. doi:10.1038/s41588-018-0085-0
- **[19]** Ting PY, et al. A molecular glue degrader of the WIZ transcription factor for fetal hemoglobin induction. *Science* 2024. doi:10.1126/science.adk6129
- **[20]** Development of a ZBTB7A and WIZ dual degrading, HbF-activating CELMoD for the treatment of sickle cell disease. *Blood* 2024; 144(Suppl 1): 169. doi:10.1182/blood-2024-203653
- **[21]** Targeted protein degradation for fetal hemoglobin induction: a new paradigm in beta-hemoglobinopathy therapy. *Biomarker Research* 2026. doi:10.1186/s40364-026-00911-5
- **[22]** BMS-986470 Phase 1/2 study in sickle cell disease. ClinicalTrials.gov NCT06481306

R.3 BCL6 ligand-directed degrader (Section 6.2)

- **[23]** Ye BH, et al. The BCL-6 proto-oncogene controls germinal-centre formation and Th2-type inflammation. *Nat Genet* 1997. doi:10.1038/ng1197-161
- **[24]** Dent AL, et al. Control of inflammation, cytokine expression, and germinal center formation by BCL-6. *Science* 1997. doi:10.1126/science.276.5312.589
- **[25]** Cattoretti G, et al. Deregulated BCL6 expression recapitulates the pathogenesis of human diffuse large B cell lymphomas in mice. *Cancer Cell* 2005. doi:10.1016/j.ccr.2005.03.037
- **[26]** Schmitz R, et al. Genetics and pathogenesis of diffuse large B-cell lymphoma. *NEJM* 2018. doi:10.1056/NEJMoa1801445
- **[27]** Chapuy B, et al. Molecular subtypes of diffuse large B cell lymphoma are associated with distinct pathogenic mechanisms and outcomes. *Nat Med* 2018. doi:10.1038/s41591-018-0016-8
- **[28]** Slabicki M, et al. Small-molecule-induced polymerization triggers degradation of BCL6. *Nature* 2020. doi:10.1038/s41586-020-2925-1
- **[29]** Discovery of BMS-986458, a potent and selective B-cell lymphoma 6 protein ligand-directed degrader, for the treatment of B-cell non-Hodgkin lymphoma. *J Med Chem* 2026. doi:10.1021/acs.jmedchem.5c03123
- **[30]** BMS-986458 Phase 1/2 study in relapsed or refractory non-Hodgkin lymphoma. ClinicalTrials.gov NCT06090539
- **[31]** Bristol Myers Squibb. Lymphoma data at ASH 2025 (press release, December 2025). news.bms.com

R.4 CD19 cell therapy (Section 6.3)

- **[32]** Mackensen A, et al. Anti-CD19 CAR T cell therapy for refractory systemic lupus erythematosus. *Nat Med* 2022. doi:10.1038/s41591-022-02017-5
- **[33]** Müller F, et al. CD19 CAR T-cell therapy in autoimmune disease: a case series with follow-up. *NEJM* 2024. doi:10.1056/NEJMoa2308917
- **[34]** van Zelm MC, et al. An antibody-deficiency syndrome due to mutations in the CD19 gene. *NEJM* 2006. doi:10.1056/NEJMoa051568
- **[35]** Bentham J, et al. Genetic association analyses implicate aberrant regulation of innate and adaptive immunity genes in the pathogenesis of systemic lupus erythematosus. *Nat Genet* 2015. doi:10.1038/ng.3434

- **[36]** Zola-cel clinical data: *Arthritis Rheumatol* 2025; 77(Suppl 9), Abstracts 0843, 1529 and LB14; *Transplant Cell Ther* 2026; 32(2): S1. Academic CD19 CAR T study adapted from Taubmann J, et al., EULAR 2023.

R.5 AI and the future state (Section 8)

- **[37]** Margolis S, et al. DrugTargetWorld: a synthetic biobank for training and benchmarking AI scientists. *arXiv* 2026. arXiv:2610.09558
- **[38]** McKinsey & Company. *AI drug discovery: focusing on what matters most* (September 2026).